Many-core and GPU developments in the parallel ELectronic Structure Infrastructure library (ELSI)

Alvaro Vazquez-Mayagoitia
alvaro@anl.gov

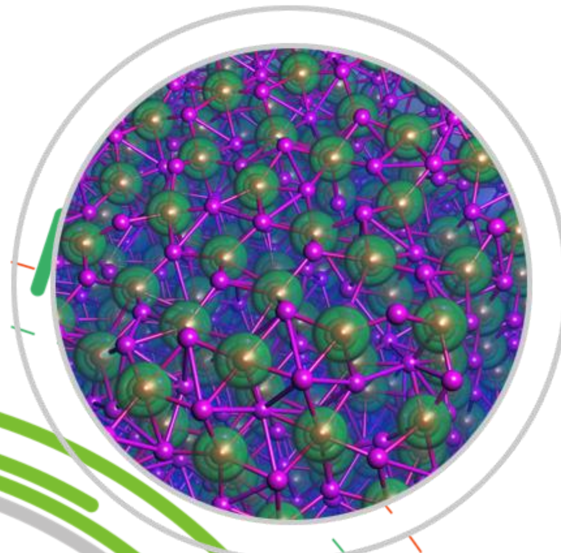

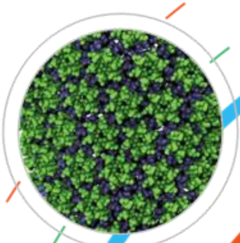

COE PP Meeting
Glendale, AZ
April 19th, 2016

Argonne Leadership Computing Facility

Argonne NATIONAL LABORATORY

ALCF Acknowledgement

This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

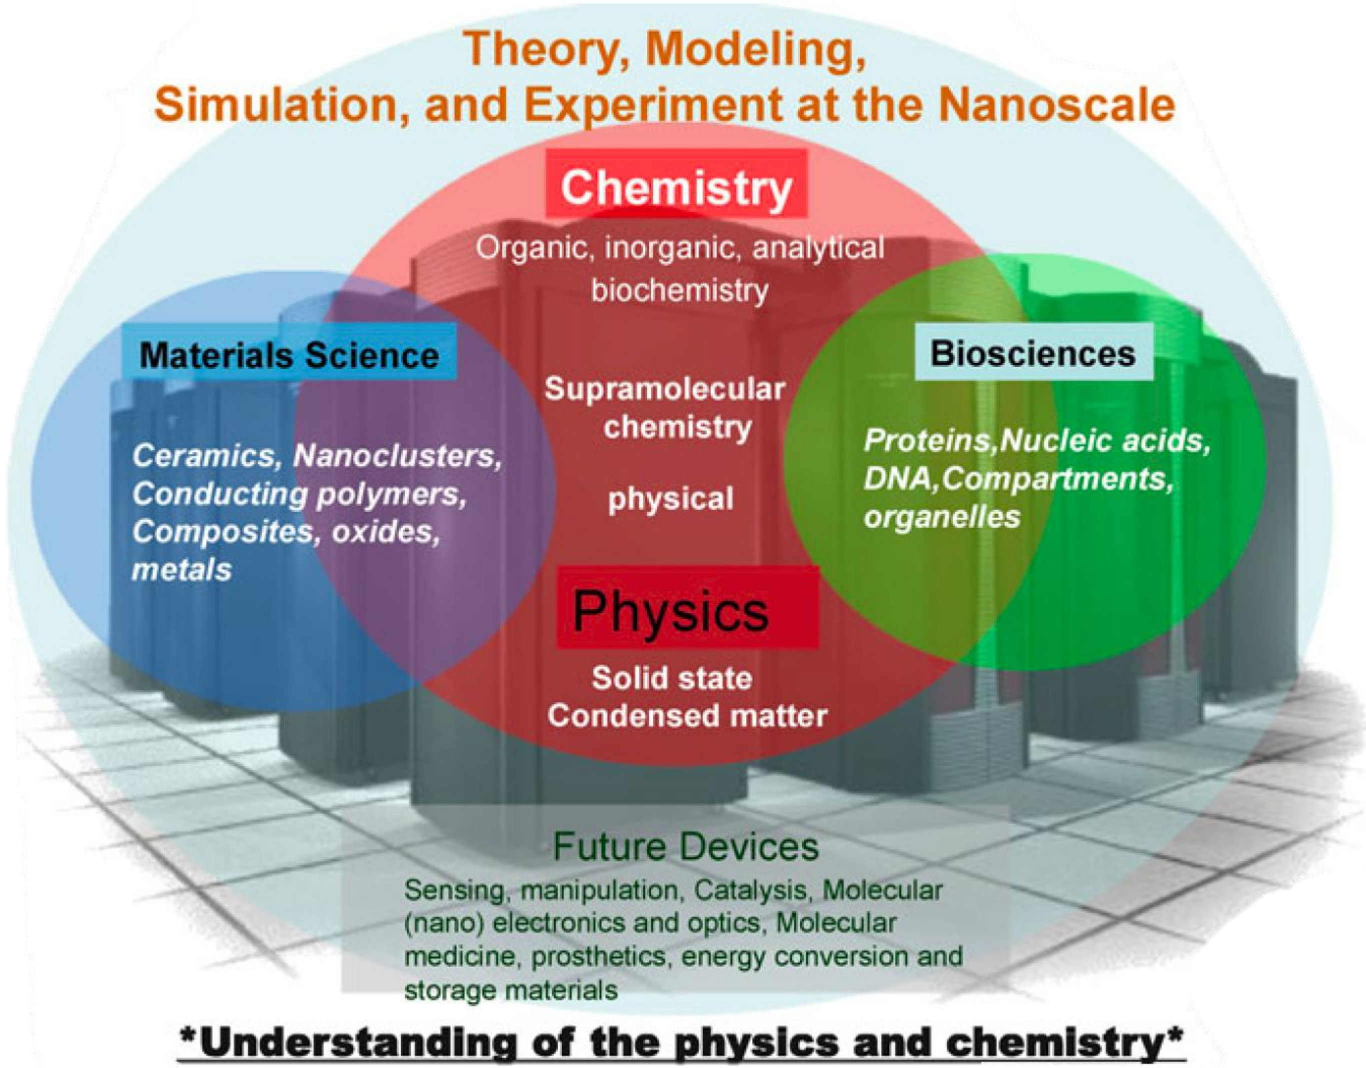

Argonne is a center of excellence for Materials Research

Usage of ASCR Facilities

2016 INCITE Allocations at ALCF

- INCITE projects are targeted to a few, very large science projects.
- >50% of the time for allocations is spent in atomics scale calculations

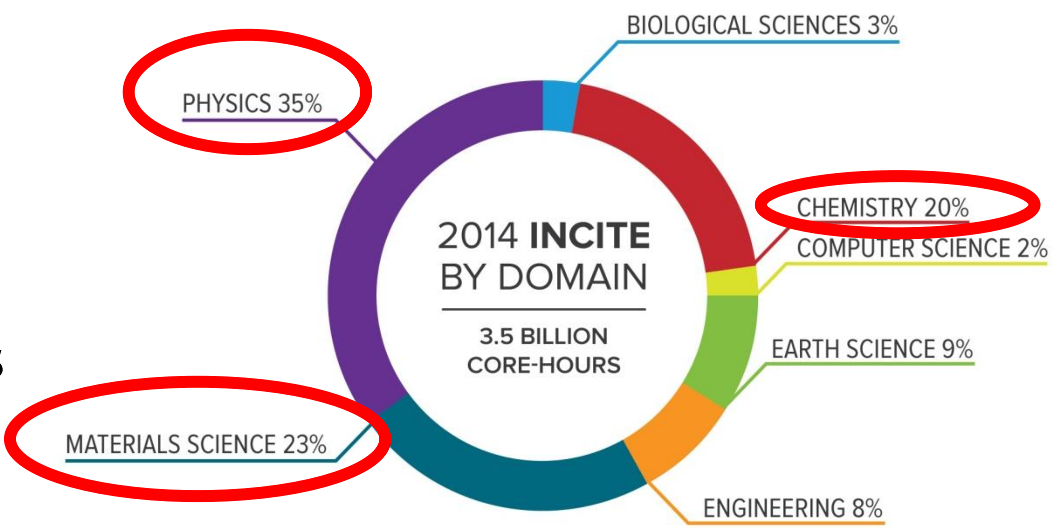

INCITE 60%, ALCC 30%, Discretionary 10%

http://science.energy.gov/ascr/highlights/

"Phase diagram" of calculation cost in computational chemistry

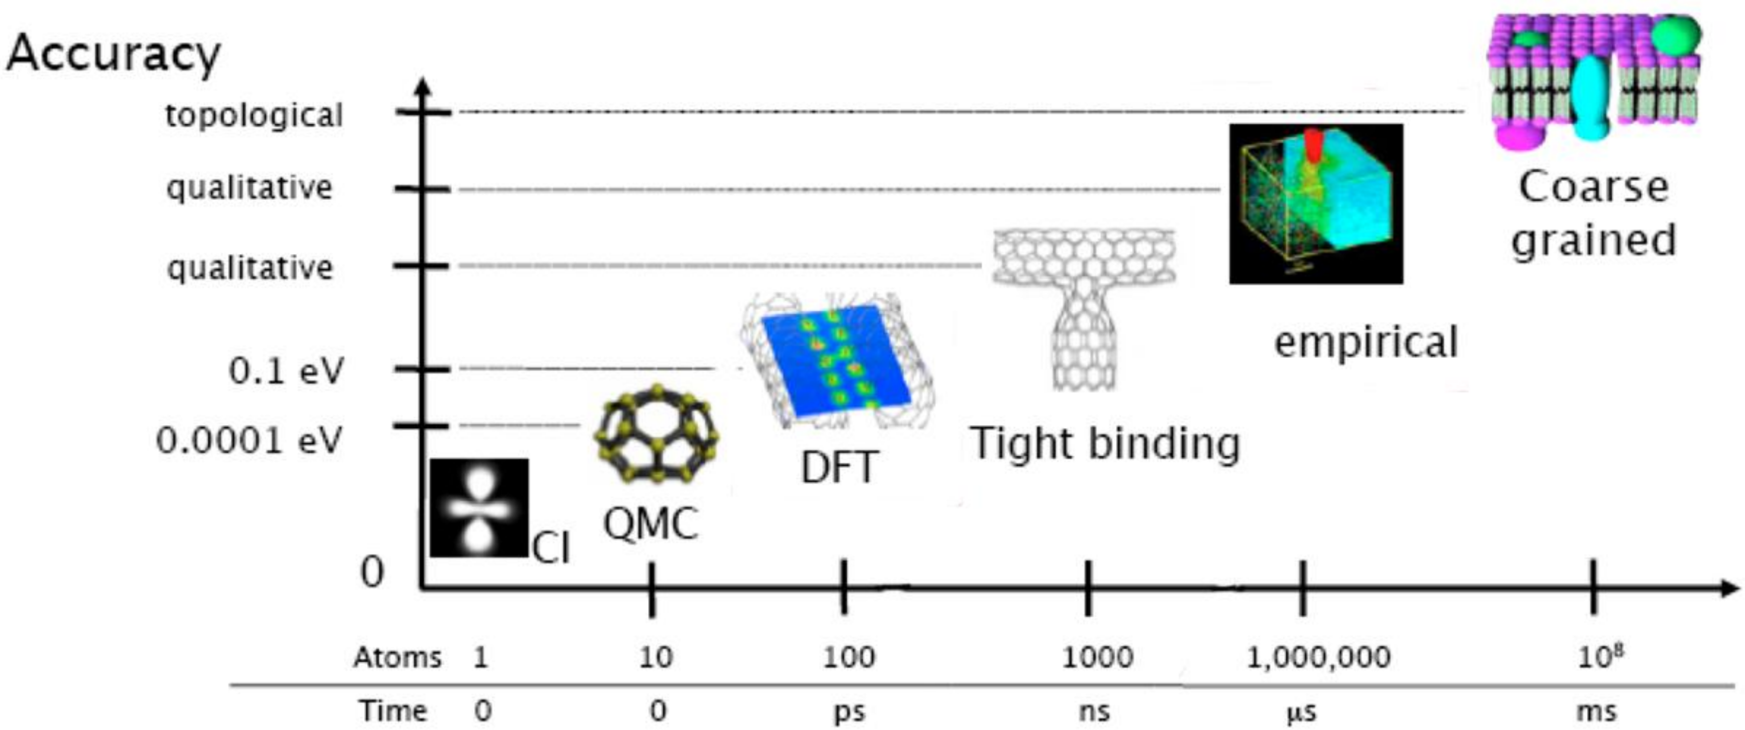

"Phase diagram" of calculation cost in computational chemistry

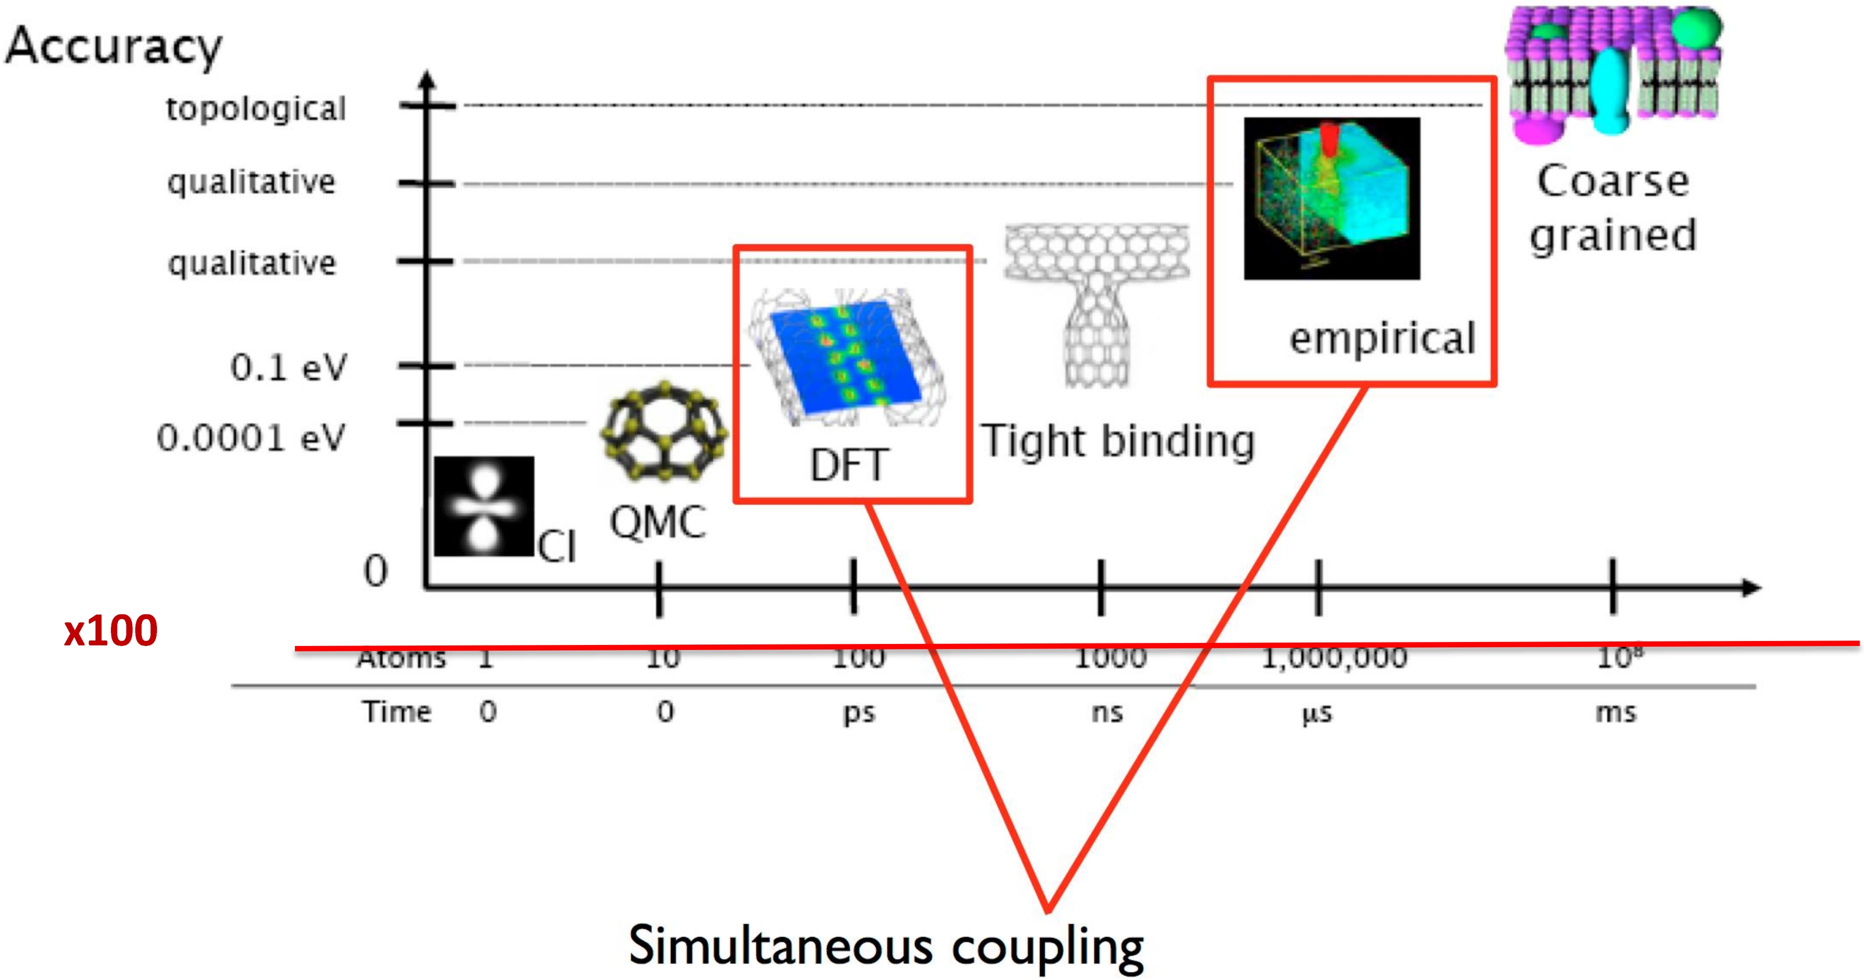

Computational Chemistry Dilemma

- Computational chemistry simulations require accurate methods at meaningful scales (space,time), which consequently need large computing resources.
- A very few software for petaflop computations for electronic structure.
- Computer power duplicates every ~ 2 year.
- New architectures challenge programmers and users for science applications.
- Long life computer programs require large community support

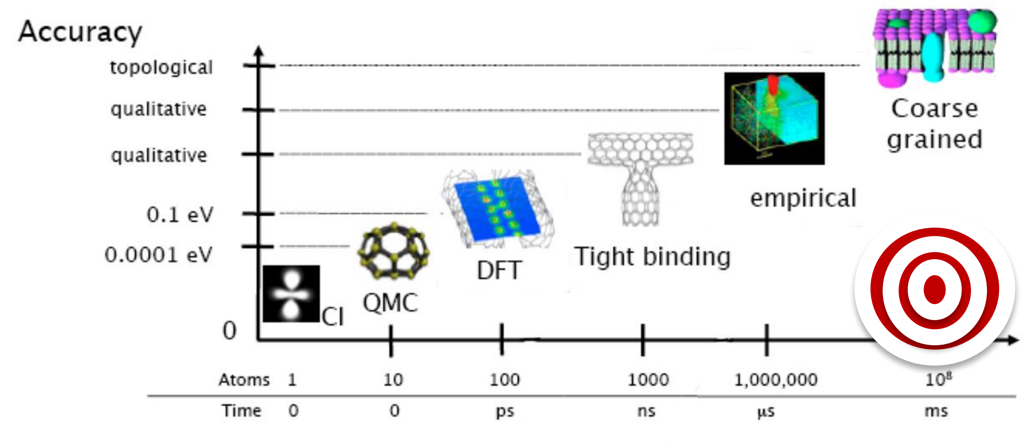

High Accuracy-
High Cost

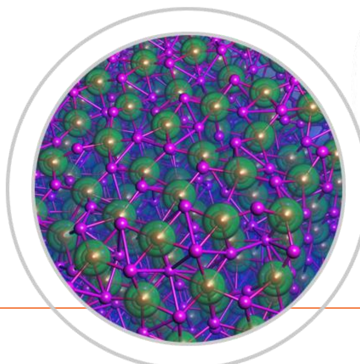

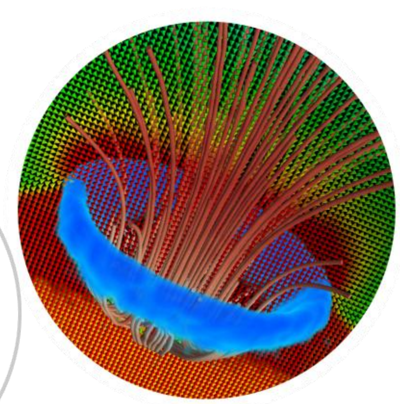

Argonne Leadership Computing Facility

Kohn-Sham scheme

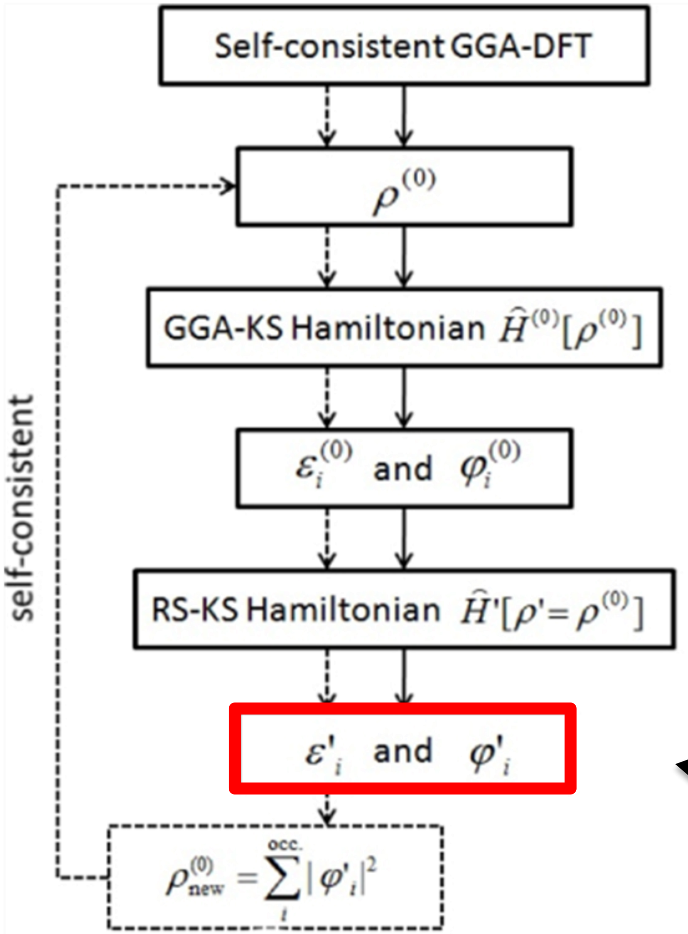

Density Functional Theory

$$E_{Total} = E[\rho(\vec{r})]$$

Auxiliary KS function

$$\Psi^{KS} = \sqrt{N!} det|\varphi_1(\vec{r_1})...\varphi_N(\vec{r_B})|$$

Density as sum of orbital densities

$$\rho(\vec{r}) = \sum_i |\varphi_i(\vec{r})|^2$$

Differential problem is now an Eigenvalue problem

2^{nd} KS theorem:
Exact ground state density minimizes *E*

Eigenvalue Solver dominates large calculations

- a real physical problems (layers graphene on SiC)

$$(6\sqrt{3} \times 6\sqrt{3}) - R30^o$$

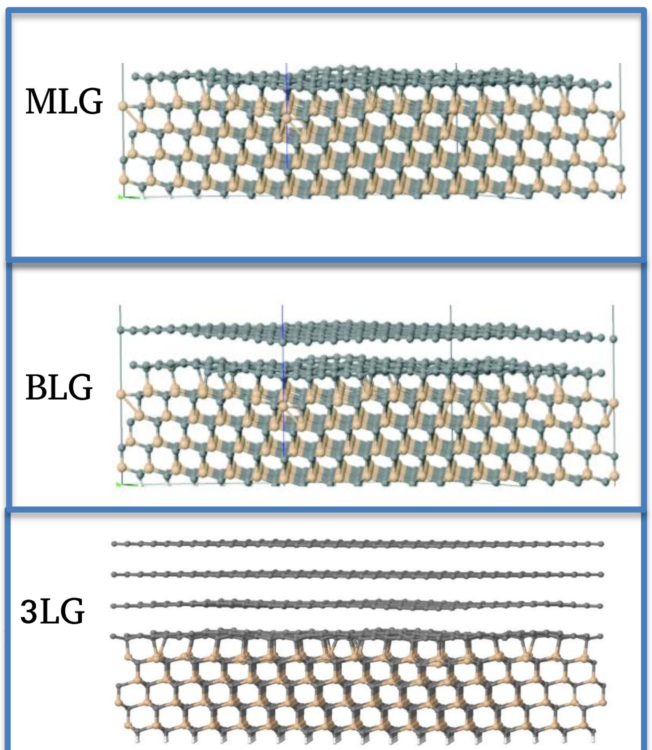

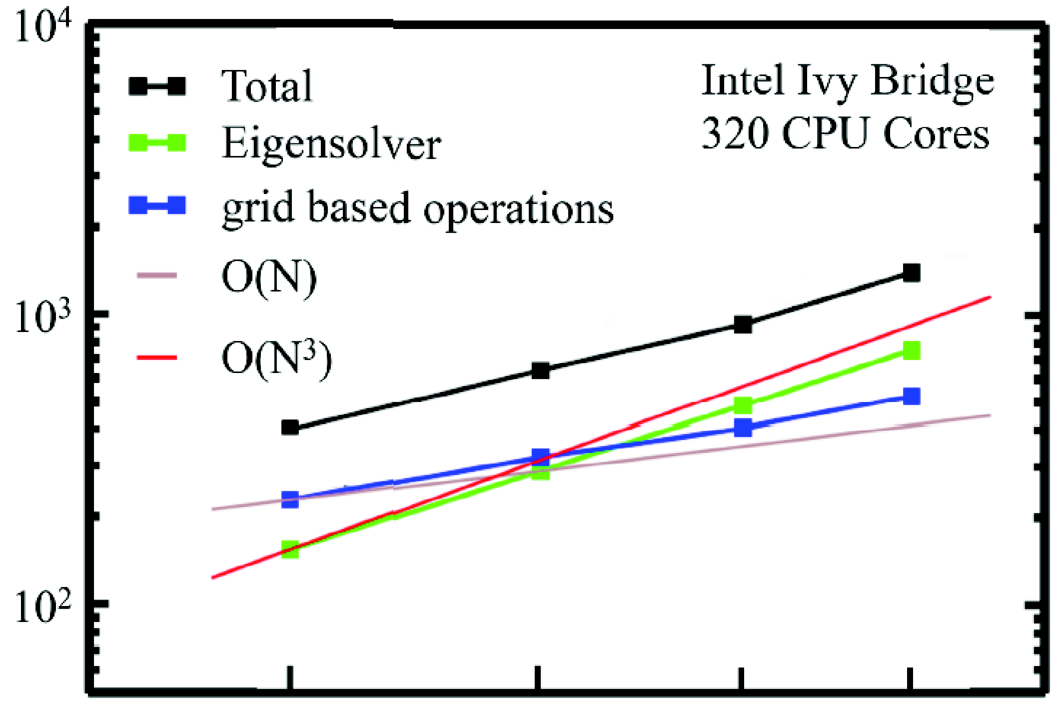

Model	"ZLG"	"MLG"	"BLG"	"3LG"
N_b	48498	61680	74862	88044
N_{at}	1310	1648	1986	2324

*for one self-consistent step

3LG:

3 layer graphene on SiC

88044 basis functions

2324 atoms

FHI-aims code

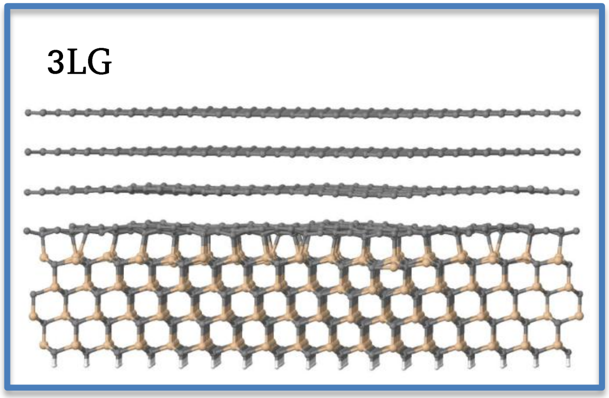

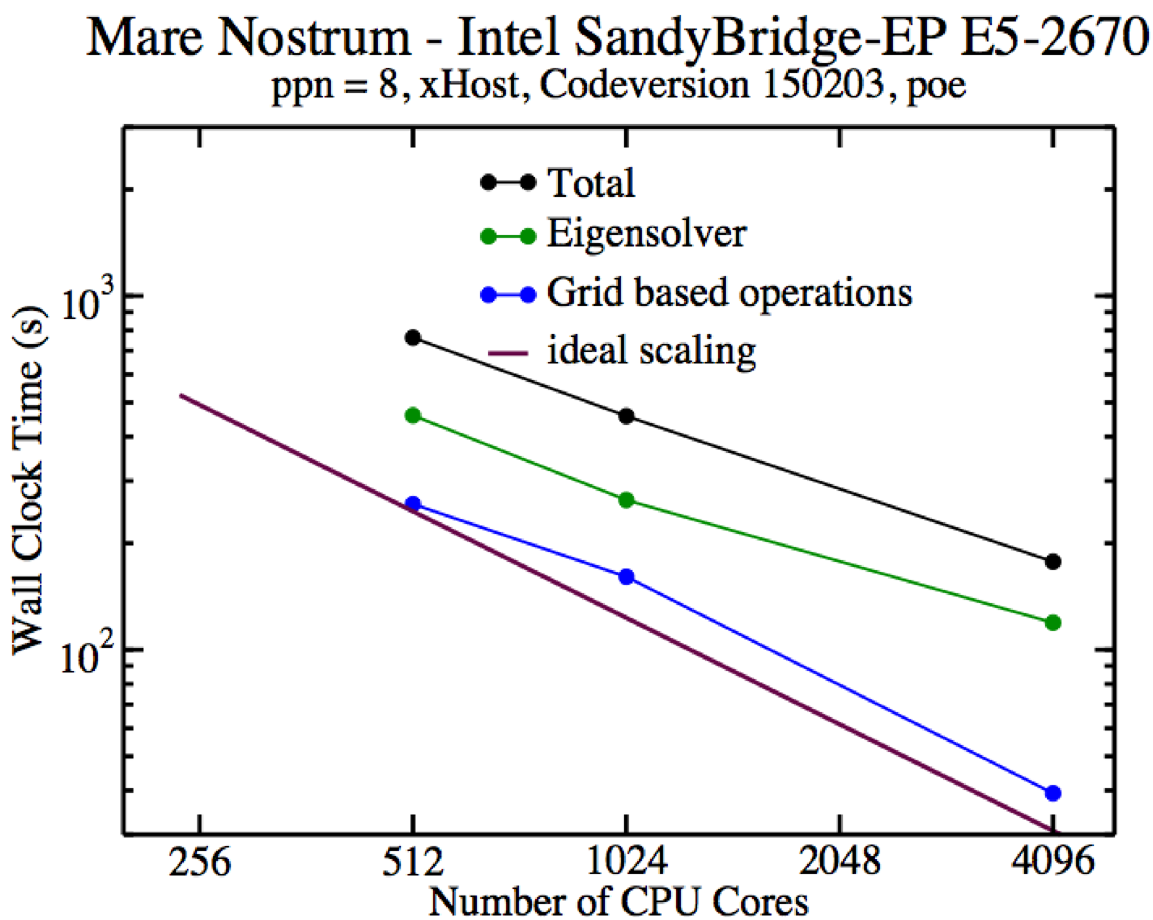

ELSI: ELectronic Structure Infrastructure

http://www.elsi-interchange.org

*Volker Blum

Volker Blum (Duke University)
Jianfeng Lu (Duke University)
Lin Lin (University of California at Berkeley)
Chao Yang (Lawrence Berkeley National Laboratory)
Alvaro Vazquez-Mayagoitia (Argonne National Laboratory)
Fabiano Corsetti (Imperial College, London)

A library to accelerate electronic structure simulations

ELSI

Electronic Structure

Community Codes

- FHI-aims
- SIESTA
- Quantum Espresso
- VASP
- AbInit
- GPAW
- Wien2k
- etc.

ELSI Interface

- **Performance**: link to high performant libraries used on modern HPC facilities
- **Ease-of-Use**: standardized interface to libraries which solve or circumvent the Kohn-Sham eigenvalue problem as well as access to their distributed storage environment
- **Range-of-Use**: widens the range of possible systems to be tackled by accessing the library suitable to the problem

Small System Sizes
ELPA
Standard $O(N^3)$
Eigenvalue Solver

Medium System Sizes
libOMM
density matrix based $O(N^3)$ library

Large System Sizes
PEXSI
density matrix based $O(N^2)$ library

**NFS funded

ELPA

Support for:

- Fast dense eigenvalue solver. Designed for high-performance architectures.
- Useful for electronic structure and many other applications. GPL license. F90 code.
- Uses a Block-cyclic distribution scheme mapping the group processors in a 2D grid. Relies on BLACS matrix parallelization layout.
- Uses efficiently external algebra libraries as BLAS, MKL, ESSL, etc.
- Main kernels in QR algorithms are hardcoded with intrinsic processor instructions (QPX, SSE, AVX, etc).
- Solves for real or complex, and for entire or partial number of eigenvalues

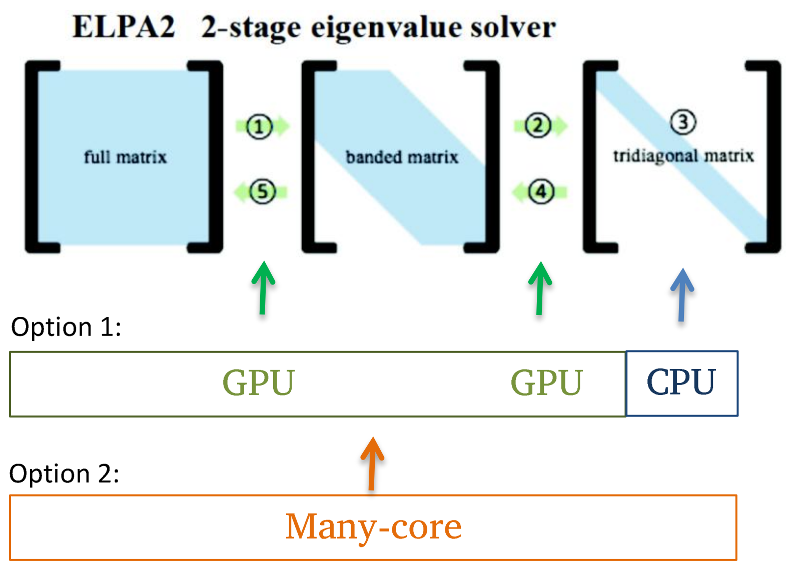

Marek et al, *The Journal of Physics: Condensed Matter* 26, 213201 (2014)

PEXSI

Electron density could be evaluated without diagonalizing Kohn-Sham-Fock Matrix

$$\rho = \mathrm{diag}(\Gamma)$$

Pole expansion

$$\Gamma \approx Im\left[\sum_{i}^{M} \omega_i (H - z_i I)^{-1}\right]$$

Selected inversion using sparse matrix techniques

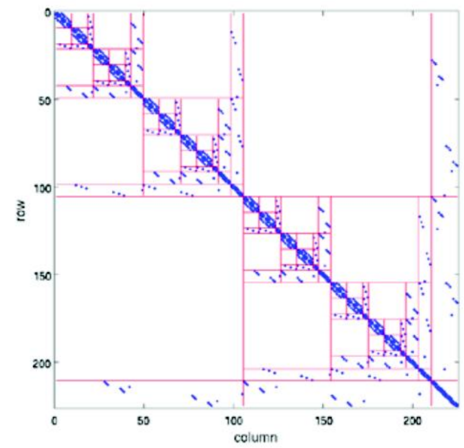

Massively parallel distributed memory implementation

- Can use 1-100K processors.
- OpenMP ready, GPU in progress.

Current support for: DGDFT, Siesta, CP2K, QuantumEspresso

- C++ code with interfaces for Fortran and C.
- Relies on SuperLU and Par/Metis.

Lin L, Lu J, Ying L and E W 2009 *Chin. Ann. Math.* **30B** 729

LibOMM

Orbital Minimization Method.- Iterative minimization algorithm originally devised for linear-scaling DFT (SIESTA code)

$$E[C] = 2tr\left\{\left[2\tilde{I} - \tilde{S}_w\right]\tilde{H}_w\right\}$$

The basic strategy: find the N/2 Wannier functions describing the occupied subspace for an N-electron system by direct unconstrained minimization

Advantages:

1. No orthogonal constrain/no explicit orthogonalization
2. Exactly quartic line search for steepest descent or conjugate gradient
3. Reuse the converged solution in the previous SCF iteration as the initial guess
4. Possible extension to a linear scaling method

F Corsetti - Computer Physics Communications, 2014

Conclusions

1. Most chemistry and physics simulations codes rely on external community supported libraries for linear algebra (Scalapack, ELPA, Elemental, etc).

2. The solution of eigen value problems is one of the most computational demanding procedure in atomic-scale calculations. For example, within the Kohn-Sham scheme, this effort could cost 60% of the total calculation time.

3. ELSI is a library that offers fast alternatives to solve the one of the major bottlenecks in Density Functional Theory which is the minimization of the energy.

4. ELSI library is the conjunction of well established algorithms with the aim to speed up electronic structure calculations in massive parallel environments.

5. Some ELSI stakeholders and partner codes have thousands of users. We aim to benefit a wide spectrum of scientists.

6. We plan to support both graphic accelerators and many-core architectures to make the most with DOE's CPU/hrs.

Thanks…